

Towards Smart Proof Search for Isabelle

Yutaka Nagashima

Data61, CSIRO, Sydney, New South Wales, Australia
first_name.last_name@data61.csiro.au

Abstract

Despite the recent progress in automatic theorem provers, proof engineers are still suffering from the lack of powerful proof automation. In this position paper we first report our proof strategy language based on a meta-tool approach. Then, we propose an AI-based approach to drastically improve proof automation for Isabelle, while identifying three major challenges we plan to address for this objective.

1 PSL and Meta-Tool Approach

In the last decade, we have seen the successful application of automated theorem provers to assist interactive theorem proving. Despite the popularity of these so-called “hammer-style” tools, their performance is still suffering from the gap between underlying logics [1].

To circumvent this problem, we introduced a proof strategy language, PSL [4], to Isabelle/HOL [5], taking a meta-tool approach. A proof strategy is an abstract description of how to attack proof obligations. Users write strategies in PSL based on their intuitions on a conjecture. Using a strategy, PSL’s runtime system generates many invocations of Isabelle’s native proof tools, called *tactics*. The runtime tries to find out the appropriate combination of tactics for each goal by trial-and-error. This way, PSL reduces the domain specific procedure of interactive theorem proving to the well-known dynamic tree search problem. The meta-tool language approach brought the following advantages:

- Domain specific procedures can be added as new sub-tools.
- Sub-tools (including Sledgehammer) can be improved independently of PSL.
- Generated efficient-proof scripts are native Isabelle proofs.

We provided a default strategy, `try_hard`. Our evaluation shows that PSL based on `try_hard` significantly outperforms Sledgehammer for many use cases [4]. However, PSL’s proof search procedure is still mostly brute-force. For example, when PSL generates many tactics, even though each of these is tailored for the proof goal utilizing the runtime information, it is still the statically written strategy, `try_hard`, that decides:

- what kind of tactics to generate,
- how to combine them, and
- in which order to apply generated tactics.

2 Meta-Tool Based Smart Proof Search

On a fundamental level, this lack of flexibility stems from the procedural nature of Isabelle’s tactic language, which describes how to prove conjectures in a step-by-step manner rather than what to do with conjectures. However, a conventional declarative language would not be a good solution either, since in many cases we cannot determine what to do with a given proof goal with certainty. Our guess may or may not be right: given a proof goal, we cannot tell which tactics to use until we complete the goal.

Probabilistic Declarative Proof Strategy Language. We propose the development of a probabilistic declarative proof strategy language (PDPSL), which allows for the description of feasible tactics for *arbitrary* proof obligation. For example, when we apply a proof strategy, **str**, written in PDPSL, to a proof search graph, **pgraph**, PDPSL’s runtime:

- (1) picks up the most promising node in **pgraph**,
- (2) applies the most promising tactic that has not been applied to that node yet, and
- (3) adds the resultant node to **pgraph**, connecting them with a labelled edge representing the tactic application and weight indicating how promising that tactic application is.

The runtime keeps applying this procedure until it reaches a solved state, in which the proof is complete. This way, PDPSL executes a best-first search at runtime.

The major challenge is the design of PDPSL. A strategy written in PDPSL is applied repeatedly during a best-first search against emerging proof obligations, and we cannot predict how these intermediate goals look like prior to the search; therefore, weighting of tactics against concrete proof goals is not good enough. PDPSL should only describe the meta-information on proof goals and information in Isabelle’s standard library. Such meta-information includes:

- Which ITP mechanism was used to define constants that appear in the proof obligation?
- Does the conjecture involve recursive functions?

This may sound too restrictive; however, when proof engineers use ITPs, we often reason on the meta level. For instance, Isabelle’s tutorial [5] introduces the following heuristics:

“Theorems about recursive functions are proved by induction.”
“The right-hand side of an equation should (in some sense) be simpler than the left-hand side.”

These are independent of user-defined formalisation.

Posterior Proof Attempt Evaluation. Evaluating a proof goal to write promising tactics is only half the story of typical interactive proof development. The other half is about evaluating the results of tactic applications. We presume that posterior evaluations on tactic application will improve proof search when combined with prior expectations on tactics described in PDPSL. The main challenge is to find a useful measure by:

- discovering (possibly multiple) useful measures that indicate progress in proof search, and
- investigating how to integrate these measures into the best-first search of PDPSL.

Of course it is not possible to come up with a heuristic that is guaranteed to work for all kinds of proof obligations. But when focusing on a particular problem domain such as systems verification, it is plausible to find a useful measure to narrow the search space at runtime.

Reinforcement Learning of PDPSL using Large Proof Corpora. PDPSL is primarily for engineers to encode their intuitions. But at the next step, we plan to improve this hand-written heuristics using reinforcement learning on the existing proof corpora. The Isabelle community has a repository of formal proofs, called the Archive of Formal Proofs (AFP) [3]. As of 2015, the size of the AFP is larger than one million lines of code. Additionally, the seL4 project open-sourced roughly 400,000 lines of Isabelle proof script [2]. We speculate that these large proof corpora will work as the basis for reinforcement learning of strategies written in PDPSL. The lack of concrete information on proof obligations can be an advantage at this stage: since PDPSL cannot describe concrete proof obligations, it is less likely to cause over-fitting.

References

- [1] Jasmin Christian Blanchette, Cezary Kaliszyk, Lawrence C. Paulson, and Josef Urban. Hammering towards QED. *J. Formalized Reasoning*, 9(1):101–148, 2016.
- [2] Gerwin Klein. Proof Engineering Challenges for Large-Scale Verification. <http://www.ai4fm.org/ai4fm-2014/>, May 2014.
- [3] Gerwin Klein, Tobias Nipkow, and Larry Paulson. The Archive of Formal Proofs. <https://www.isa-afp.org/>, 2004–2016.
- [4] Yutaka Nagashima and Ramana Kumar. A Proof Strategy Language and Proof Script Generation for Isabelle. *CoRR*, abs/1606.02941, 2016.
- [5] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer, 2002.